Balanced Allocation on Graphs: A Random Walk Approach

Ali Pourmiri Institute for Research in Fundamental Sciences (IPM), Tehran, Iran
pourmiri@ipm.ir

March 1, 2016

arXiv:1407.2575v4 [cs.DS] 27 Feb 2016

Abstract

The standard balls-into-bins model is a process which randomly allocates m balls into n bins where each ball picks d bins independently and uniformly at random and the ball is then allocated in a least loaded bin in the set of d choices. When $m = n$ and $d = 1$, it is well known that at the end of process the maximum number of balls at any bin, the *maximum load*, is $(1 + o(1))\frac{\log n}{\log \log n}$ with high probability[1]. Azer et al. [4] showed that for the d-choice process, $d \geqslant 2$, provided ties are broken randomly, the maximum load is $\frac{\log \log n}{\log d} + \mathcal{O}(1)$.

In this paper we propose algorithms for allocating n sequential balls into n bins that are interconnected as a d-regular n-vertex graph G, where $d \geqslant 3$ can be any integer. Let l be a given positive integer. In each round t, $1 \leqslant t \leqslant n$, ball t picks a node of G uniformly at random and performs a non-backtracking random walk of length l from the chosen node. Then it allocates itself on one of the visited nodes with minimum load (ties are broken uniformly at random). Suppose that G has a sufficiently large girth and $d = \omega(\log n)$. Then we establish an upper bound for the maximum number of balls at any bin after allocating n balls by the algorithm, called *maximum load*, in terms of l with high probability. We also show that the upper bound is at most an $\mathcal{O}(\log \log n)$ factor above the lower bound that is proved for the algorithm. In particular, we show that if we set $l = \lfloor (\log n)^{\frac{1+\epsilon}{2}} \rfloor$, for every constant $\epsilon \in (0, 1)$, and G has girth at least $\omega(l)$, then the maximum load attained by the algorithm is bounded by $\mathcal{O}(1/\epsilon)$ with high probability. Finally, we slightly modify the algorithm to have similar results for balanced allocation on d-regular graph with $d \in [3, \mathcal{O}(\log n)]$ and sufficiently large girth.

1 Introduction

The standard balls-into-bins model is a process which randomly allocates m balls into n bins where each ball picks d bins independently and uniformly at random and the ball is then allocated in a least loaded bin in the set of d choices. When $m = n$ and $d = 1$, it is well known that at the end of process the maximum number of balls at any bin, the *maximum load*, is $(1 + o(1))\frac{\log n}{\log \log n}$ with high probability. Azer et al. [4] showed that for the d-choice process, $d \geqslant 2$, provided ties are broken randomly, the maximum load is $\frac{\log \log n}{\log d} + \mathcal{O}(1)$. For a complete survey on the standard balls-into-bins process we refer the reader to [13]. Many subsequent works consider the settings where the choice of bins are not necessarily independent and uniform. For instance, Vöcking [15] proposed an algorithm called *always-go-left* that uses exponentially smaller number of choices and achieve a maximum load of $\frac{\log \log n}{d\phi_d} + \mathcal{O}(1)$ whp, where $1 \leqslant \phi_d \leqslant 2$ is an specified constant. In this algorithm, the bins are partitioned into d groups of size n/d and each ball picks one random bin from each group. The ball is then allocated in a least loaded bin among the chosen bins and ties are broken asymmetrically.

In many applications selecting any random set of choices is costly. For example, in peer-to-peer or cloud-based systems balls (jobs, items,...) and bins (servers, processors,...) are randomly placed in a metric space (e.g., $\mathbb{R}^2$) and the balls have to be allocated on bins that are close to them as it minimizes the access latencies. With regard to such applications, Byer et al. [7] studied a model, where n bins (servers) are uniformly at random placed on a geometric space. Then each ball in turn picks d locations in the space and allocates itself on a nearest neighboring bin with minimum load among other d bins. In this scenario, the probability that a location close to a server is chosen depends on the distribution of other servers in the space and hence there is no a uniform distribution over the potential choices. Here, the authors

[1]With high probability refers to an event that holds with probability $1 - 1/n^c$, where c is a constant. For simplicity, we sometimes abbreviate it as whp.

showed the maximum load is $\frac{\log\log n}{\log d} + \mathcal{O}(1)$ whp. Later on, Kenthapadi and Panigrahy [11] proposed a model in which bins are interconnected as a Δ-regular graph and each ball picks a random edge of the graph. It is then placed at one of its endpoints with smaller load. This allocation algorithm results in a maximum load of $\log\log n + \mathcal{O}\left(\frac{\log n}{\log(\Delta/\log^4 n)}\right) + \mathcal{O}(1)$. Peres et al. [14] also considered a similar model where number of balls m can be much larger than n (i.e., $m \gg n$) and the graph is not necessarily regular. Then, they established upper bound $\mathcal{O}(\log n/\sigma)$ for the gap between the maximum and the minimum loaded bin after allocating m balls, where σ is the edge expansion of the graph. Following the study of balls-into-bins with correlated choices, Godfrey [10] generalized the model introduced by Kenthapadi and Panigrahy such that each ball picks an random edge of a hypergraph that has $\Omega(\log n)$ bins and satisfies some mild conditions. Then he showed that the maximum load is a constant whp.

Recently, Bogdan et al. [6] studied a model where each ball picks a random node and performs a local search from the node to find a node with local minimum load, where it is finally placed on. They showed that when the graph is a constant degree expander, the local search guarantees a maximum load of $\Theta(\log\log n)$ whp.

Our Results. In this paper, we study balls-into-bins models, where each ball chooses a set of related bins. We propose allocation algorithms for allocating n sequential balls into n bins that are organized as a d-regular n-vertex graph G. Let l be a given positive integer. A non-backtracking random walk (NBRW) W of length l started from a node is a random walk in l steps so that in each step the walker picks a neighbor uniformly at random and moves to that neighbor with an additional property that the walker never traverses an edge twice in a row. Further information about NBRWs can be found in [1] and [2]. Our allocation algorithm, denoted by $\mathcal{A}(G, l)$, is based on a random sampling of bins from the neighborhood of a given node in G by a NBRW from the node. The algorithm proceeds as follows: In each round t, $1 \leqslant t \leqslant n$, ball t picks a node of G uniformly at random and performs a NBRW $W = (u_0, u_1 \ldots, u_l)$, called *$l$-walk*. After that the ball allocates itself on one of the visited nodes with minimum load and ties are broken randomly. Our result concerns bounding the maximum load attained by $\mathcal{A}(G, l)$, denoted by m^*, in terms of l. Note that if the balls are allowed to take NBRWs of length $l = \Omega(\log n)$ on a graph with girth at least l, then the visited nodes by each ball generates a random hyperedge of size $l + 1$. Then applying the Godfrey's result [10] implies a constant maximum load whp. So, for the rest of the paper we focus on NBRWs of sub-logarithmic length (i.e., $l = o(\log_d n)$). We also assume that $l = \omega(1)$ and G is a d-regular n-vertex graph with girth at least $\omega(l \log\log n)$ and $d = \omega(\log n)$. However, when $l = \lfloor(\log n)^{\frac{1+\epsilon}{2}}\rfloor$, for any constant $\epsilon \in (0, 1)$, G with girth at least $\omega(l)$ suffices as well. It is worth mentioning that there exist several explicit families of n-vertex d-regular graph with arbitrary degree $d \geqslant 3$ and girth $\Omega(\log_d n)$ (see e.g. [9]).

In order to present the upper bound, we consider two cases:

I. If $l \geqslant 4\gamma_G$, where $\gamma_G = \sqrt{\log_d n}$, then we show that whp,

$$m^* = \mathcal{O}\left(\frac{\log\log n}{\log(l/\gamma_G)}\right).$$

Thus, for a given G satisfying the girth condition, if we set $l = \lfloor(\log_d n)^{\frac{1+\epsilon}{2}}\rfloor$, for any constant $\epsilon \in (0, 1)$, then we have $l/\gamma_G \geqslant (\log n)^{\epsilon/2}$ and by applying the above upper bound we have $m^* = \mathcal{O}(1/\epsilon)$ whp.

II. If $\omega(1) \leqslant l \leqslant 4 \cdot \gamma_G$, then we show that whp,

$$m^* = \mathcal{O}\left(\frac{\log_d n \cdot \log\log n}{l^2}\right).$$

In addition to the upper bound, we prove that whp,

$$m^* = \Omega\left(\frac{\log_d n}{l^2}\right)$$

(for a proof see Appendix F). If G is a d-regular graph with $d \in [3, \mathcal{O}(\log n)]$, then we slightly modify allocation algorithm $\mathcal{A}(G, l)$ and show the similar results for m^* in l. The algorithm $\mathcal{A}'(G, l)$ for sparse graphs proceeds as follows: Let us first define parameter

$$r_G = \lceil 2 \cdot \log_{d-1} \log n \rceil.$$

For each ball t, the ball takes a NBRW of size $l \cdot r_G$, say $(u_0, u_1, \cdots, u_{lr_G})$, and then a subset of visited nodes, $\{u_{j \cdot r_G} \mid 0 \leqslant j \leqslant l\}$, called *potential choices*, is selected and finally the ball

is allocated on a least-loaded node of potential choices (ties are broken randomly). Provided G has sufficiently large girth, we show the similar upper and lower bounds as the allocation algorithm $\mathcal{A}(G, l)$ on d-regular graphs with $d = \omega(\log n)$ (see Appendix E).

Comparison with Related Works. The setting of our work is closely related to [6]. In this paper in each step a ball picks a node of a graph uniformly at random and performs a local search to find a node with local minimum load and finally allocates itself on it. They showed that with high probability the local search on expander graphs obtains a maximum load of $\Theta(\log \log n)$. In comparison to the mentioned result, our new protocol achieves a further reduction in the maximum load, while still allocating a ball close to its origin. Our result suggests a trade off between allocation time and maximum load. In fact we show a constant upper bound for sufficient long walks (i.e., $l = (\log n)^{\frac{1+\epsilon}{2}}$, for any constant $\epsilon \in (0, 1)$). Our work can also be related to the one by Kenthapadi and Panigrahy where each ball picks a random edge from a $n^{\Omega(1/\log \log n)}$-regular graph and places itself on one of the endpoints of the edge with smaller load. This model results into a maximum load of $\Theta(\log \log n)$. Godfrey [10] considered balanced allocation on hypergraphs where balls choose a random edge e of a hypergraph satisfying some conditions, that is, first the size of each edge s is $\Omega(\log n)$ and $\mathbf{Pr}[u \in e] = \Theta(\frac{s}{n})$ for any bin u. The latter one is called *balanced condition*. Berenbrink et al. [5] simplified Godfrey's proof and slightly weakened the balanced condition but since both analysis apply a Chernoff bound, it seems unlikely that one can extend the analysis for hyperedges of size $o(\log n)$. Our model can also be viewed as a balanced allocation on hypergraphs, because every l-walk is a random hyperedge of size $l + 1$ that also satisfies the balanced condition (see Lemma A.4). By setting the right parameter for $l = o(\log n)$, we show that the algorithm achieves a constant maximum load with sub-logarithmic number of choices.

In a different context, Alon and Lubetzky [2] showed that if a particle starts a NBRW of length n on n-vertex graph with high-girth then the number of visits to nodes has a Poisson distribution. In particular they showed that the maximum visit to a node is at most $(1 + o(1)) \cdot \frac{\log n}{\log \log n}$. Our result can be also seen as an application of the mathematical concept of NBRWs to task allocation in distributed networks.

Techniques. To derive a lower bound for the maximum load we first show that whp there is a path of length l which is traversed by at least $\Omega(\log_d n/l)$ balls. Also, each path contains $l+1$ choices and hence, by pigeonhole principle there is a node with load at least $\Omega\left(\log_d n/l^2\right)$, which is a lower bound for m^*. We establish the upper bound based on *witness graph* techniques. In our model, the potential choices for each ball are highly correlated, so the technique for building the witness graph is somewhat different from the one for standard balls-into-bins. Here we propose a new approach for constructing the witness graph. We also show a key property of the algorithm, called (α, n_1)-*uniformity*, that is useful for our proof technique. We say an allocation algorithm is (α, n_1)-uniform if the probability that, for every $1 \leqslant t \leqslant n_1$, ball t is placed on an arbitrary node is bounded by α/n, where $n_1 = \Theta(n)$ and $\alpha = \mathcal{O}(1)$. Using this property we conclude that for a given set of nodes of size $\Omega(\log n)$, after allocating n_1 balls, the average load of nodes in the set is some constant whp. Using witness graph method we show that if there is a node with load larger than some threshold then there is a collection of nodes of size $\Omega(\log n)$ where each of them has load larger than some specified constant. Putting these together implies that after allocating n_1 balls the maximum load, say m_1^*, is bounded as required whp. To derive an upper bound for the maximum load after allocation n balls, we divide the allocation process into n/n_1 phases and show that the maximum load at the end of each phase increases by at most m_1^* and hence $m^* \leqslant (n/n_1)m_1^*$ whp.

Discussion and Open Problems. In this paper, we proposed balls-into-bins model, where each ball picks a set of nodes that are visited by a NBRW of length l and place itself on a visited node with minimum load. One may ask whether it is possible to replace a NBRW of length l by several parallel random walks of shorter length (started from the same node) and get the similar results?

In our result we constantly use the assumption that the graph locally looks like a d-ary tree. It is also known that cycles in random regular graph are restively far from each other (e.g, see [8]), so we believe that our approach can be extended for balanced allocation on random regular graphs.

Many works in this area (see e.g.[6, 11]) assumed that the underlying networks is regular, it would be interesting to investigate random walk-based algorithms for irregular graphs.

Outline. In Section 2, we present notations and some preliminary results that are required for the analysis of the algorithm. In Section 3 we show how to construct a witness graph and

then in Section 4 by applying the results we the upper bound for the maximum load.

2 Notations, Definitions and Preliminaries

In this section we provide notations, definitions and some preliminary results. A *non-backtracking random walk* (NBRW) W of length l started from a node is a simple random walk in l steps so that in each step the walker picks a neighbor uniformly at random and moves to that neighbor with an additional property that the walker never traverses an edge twice in a row. Throughout this paper we assume that $l \in [\omega(1), o(\log_d n)]$ is a given parameter and G is a d-regular graph with girth $10 \cdot l \cdot \log\log n$. Note that we will see that the condition on the girth can be relaxed to $\omega(l)$, for any l higher than $(\log_d n)^{\frac{1+\epsilon}{2}}$, where $\epsilon \in (0,1)$ is a constant.

It is easy to see that the visited nodes by a non-backtracking walk of length l on G induces a path of length l, which is called an *l-walk*. For simplicity, we use W to denote both the l-walk and the set of visited nodes by the l-walk. Also, we define $f(W)$ to be the number of balls in a least-loaded node of W. The *height* of a ball allocated on a node is the number balls that are placed on the node before the ball.

For every two nodes $u, v \in V(G)$, let $d(u,v)$ denote the length of shortest path between u and v in G. Since G has girth at least $\omega(l)$, every path of length at most l is specified by its endpoints, say u and v. So we denote the path by interval $[u,v]$. Note that for any graph H, $V(H)$ denotes the vertex set of H.

Definition 1 (Interference Graph)**.** *For every given pair (G,l), the interference graph $\mathcal{I}(G,l)$ is defined as follows: The vertex set of $\mathcal{I}(G,l)$ is the set of all l-walks in G and two vertices W and W' of $\mathcal{I}(G,l)$ are connected if and only if $W \cap W' \neq \emptyset$. Note that if pair (G,l) is clear from the context, then the interference graph is denoted by $\mathcal{I}$.*

Now, let us interpret allocation process $\mathcal{A}(G,l)$ as follows:

For every ball $1 \leqslant t \leqslant n$, the algorithm picks a vertex of $\mathcal{I}(G,l)$, say W_t, uniformly at random and then allocates ball t on a least-loaded node of W_t (ties are broken randomly). Let $1 \leqslant n_1 \leqslant n$ be a given integer and assume that $\mathcal{A}(G,l)$ has allocated balls until the n_1-th ball. We then define $\mathcal{H}_{n_1}(G,l)$ to be the induced subgraph of $\mathcal{I}(G,l)$ by $\{W_t : 1 \leqslant t \leqslant n_1\} \subset V(\mathcal{I})$.

Definition 2. *Let λ and μ be given positive integers. We say rooted tree $T \subset \mathcal{I}(G,l)$ is a (λ,μ)-tree if T satisfies:*

1) $|V(T)| = \lambda$,

2) $|\cup_{W \in V(T)} W| \geqslant \mu$.

Note that the latter condition is well-defined because every vertex of T is an $(l+1)$-element subset of $V(G)$. A (λ,μ)-tree T is called c-loaded, if T is contained in $\mathcal{H}_{n_1}(G,l)$, for some $1 \leqslant n_1 \leqslant n$, and every node in $\cup_{W \in V(T)} W$ has load at least c.

2.1 Appearance Probability of a c-Loaded (λ,μ)-Tree

In this subsection we formally define the notion of (α, n_1)-uniformity for allocation algorithms, and then present our key lemma concerning the uniformity of $\mathcal{A}(G,l)$. By using this lemma we establish an upper bound for the probability that a c-loaded (λ,μ)-tree contained in $\mathcal{H}_{n_1}$ exists. The proof of the following lemmas can be found in Appendix A.

Definition 3. *Suppose that $\mathcal{B}$ be an algorithm that allocates n sequential balls into n bins. Then we say $\mathcal{B}$ is (α, n_1)-uniform if, for every $1 \leqslant t \leqslant n_1$ and every bin u, after allocating t balls we have that*

$$\mathbf{Pr}\left[\textit{ball } t+1 \textit{ is allocated on } u\right] \leqslant \frac{\alpha}{n},$$

where α is some constant and $n_1 = \Theta(n)$.

Lemma 2.1 (Key Lemma)**.** *$\mathcal{A}(G,l)$ is an (α, n_1)-uniform allocation algorithm, where $n_1 = \lfloor n/(6\mathrm{e}\alpha) \rfloor$.*

In the next lemma, we derive an upper bound for the appearance probability a c-loaded (λ,μ)-tree, whose proof is inspired by [11, Lemma 2.1].

Lemma 2.2. *Let λ, μ and c be positive integers. Then the probability that there exists a c-loaded (λ,μ)-tree contained in $\mathcal{H}_{n_1}(G,l)$ is at most*

$$n \cdot \exp(4\lambda \log(l+1) - c\mu).$$

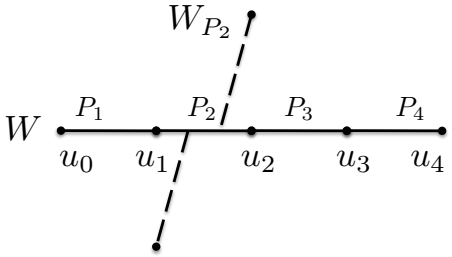

Figure 1: The Partition step on W for $k = 4$ and the Branch step for P_2 that gives W_{P_2}, shown by dashed line.

3 Witness Graph

In this section, we show that if there is a node whose load is larger than a threshold, then we can construct a c-loaded (λ, μ)-tree contained in $\mathcal{H}_{n_1}(G, l)$. Our construction is based on an iterative application of a 2-step procedure, called Partition-Branch. Before we explain the construction, we draw the reader's attention to the following remark:

Remark. The intersection (union) of two arbitrary graphs is a graph whose vertex set and edge set are the intersection (union) of the vertex and edge sets of those graphs. Let $\cap_g$ and $\cup_g$ denote the graphical intersection and union. Note that we use $\cap$ ($\cup$) to denote the set intersection (union) operation. Moreover, since G has girth $\omega(l)$, the graphical intersection of every two l-walks in G is either empty or a path (of length $\leqslant l$). Recall that W denotes both an l-walk and the set of nodes in the l-walk.

Partition-Branch. Let $k \geqslant 1$ and $\rho \geqslant 1$ be given integers and W be an l-walk with $f(W) \geqslant \rho + 1$. The Partition-Branch procedure on W with parameters ρ and k, denoted by $PB(\rho, k)$, proceeds as follows:

Partition: It partitions W into k edge-disjoint subpaths:

$$\mathcal{P}_k(W) = \{[u_i, u_{i+1}] \subset W, 0 \leqslant i \leqslant k-1\},$$

where $d(u_i, u_{i+1}) \in \{\lfloor l/k \rfloor, \lceil l/k \rceil\}$.

Branch: For a given $P_i = [u_i, u_{i+1}] \in \mathcal{P}_k(W)$, it finds (if exists) another l-walk W_{P_i} intersecting P_i that satisfies the following conditions:

(C1) $\emptyset \neq W_{P_i} \cap W \subseteq P_i \setminus \{u_i, u_{i+1}\}$.

(C2) $f(W_{P_i}) \geqslant f(W) - \rho$.

We say procedure $PB(\rho, k)$ on a given l-walk W is *valid*, if for every $P \in \mathcal{P}_k(W)$, W_P exists. We usually refer to W as the father of W_P. For a graphical view of the Partition-Branch procedure see Figure 1.

Definition 4 (Event $\mathcal{N}_\delta$)**.** *For any given $1 \leqslant \delta \leqslant l$, we say that event $\mathcal{N}_\delta$ holds, if after allocating at most n balls by $\mathcal{A}(G, l)$, every path of length δ is contained in less than $6 \log_{d-1} n/\delta$ l-walks that are randomly chosen by $\mathcal{A}(G, l)$.*

For the sake of construction, let us define a set of parameters, depending on d, n, and l, which are used throughout the paper

$$\begin{aligned} k &:= \max\{4, \lfloor l/\sqrt{\log_d n} \rfloor\}, \\ \delta &:= \lfloor \lfloor l/k \rfloor / 4 \rfloor, \\ \rho &:= \lceil 6 \log_d n/\delta^2 \rceil. \end{aligned}$$

Lemma 3.1. *Suppose that event $\mathcal{N}_\delta$ holds and W be an l-walk with*

$$f(W) \geqslant \rho + 1.$$

Then the procedure $PB(\rho, k)$ on W is valid.

For a proof see Appendix B.

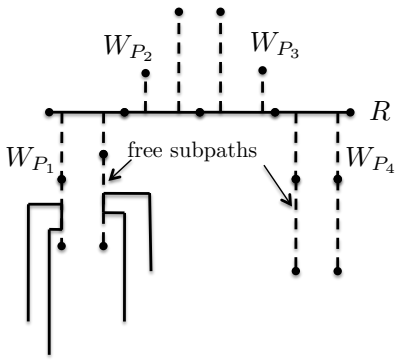

Figure 2: The first level $\mathcal{L}_1 = \{W_{P_1}, W_{P_2}, W_{P_3}, W_{P_4}\}$ and the Branch step for free subpaths of $\mathcal{P}_k(W_{P_1})$.

3.1 Construction of Witness Graph

In this subsection, we show how to construct a c-loaded (λ, μ)-tree contained in $\mathcal{H}_{n_1}$. Let $U_{n_1,l,h}$ denote the event that after allocating at most $n_1 \leqslant n$ balls by $\mathcal{A}(G, l)$ there is a node with load at least $h\rho + c + 1$, where $c = \mathcal{O}(1)$ and $h = \mathcal{O}(\log \log n)$ are positive integers that will be fixed later. Suppose that event $U_{n_1,l,h}$ conditioning on $\mathcal{N}_\delta$ happens. Then there is an l-walk R, called root, that corresponds to the ball at height $h\rho + c$ and has $f(R) \geqslant h\rho + c$. Applying Lemma 3.1 shows that $PB(\rho, k)$ on R is valid. So, let us define

$$\mathcal{L}_1 := \{W_P, P \in \mathcal{P}_k(R)\},$$

which is called the first level and R is the father of all l-walks in $\mathcal{L}_1$. (C2) in the Partition-Branch procedure ensures that for every $W \in \mathcal{L}_1$,

$$f(W) \geqslant (h-1)\rho + c.$$

Once we have the first level we recursively build the i-th level from the $(i-1)$-th level, for every $2 \leqslant i \leqslant h$. We know that each W except R is created by the Branch step on its father. Let us fix $W \in \mathcal{L}_{i-1}$ and its father W'. We then apply the Partition step on W and get $\mathcal{P}_k(W)$. We say $P \in \mathcal{P}_k(W)$ is a *free* subpath if it does not share any node with W'. By (C1), we have that $\emptyset \neq W \cap W' = [u, v] \subset P'$, for some $P' \in \mathcal{P}_k(W')$ and hence $d(u, v) \leqslant \lceil l/k \rceil$. So, $[u, v]$ shares node(s) with at most 2 subpaths in $\mathcal{P}_k(W)$ and thus $\mathcal{P}_k(W)$ contains at least $k-2$ free subpaths. Let $\mathcal{P}_k^0(W) \subset \mathcal{P}_k(W)$ denote an arbitrary set of free subpaths of size $k-2$. By (C2) and the recursive construction, we have that $f(W) \geqslant (h-i+1)\rho + c$, for each $W \in \mathcal{L}_{i-1}$. Therefore, by Lemma 3.1, $PB(\rho, k)$ on W is valid. Now we define the i-th level as follows,

$$\mathcal{L}_i = \bigcup_{W \in \mathcal{L}_{i-1}} \{W_P, P \in \mathcal{P}_k^0(W)\}.$$

For a graphical view see Figure 2. The following lemma guarantees that our construction gives a c-loaded (λ, μ)-tree in $\mathcal{H}_{n_1}$ with desired parameters (for a proof see Appendix C).

Lemma 3.2. *Suppose that G has girth at least $10hl$ and $U_{n_1,l,h}$ conditioning on $\mathcal{N}_\delta$ happens. Then there exists a c-loaded (λ, μ)-tree $T \subset \mathcal{H}_{n_1}$, where $\lambda = 1 + k\sum_{j=0}^{h-1}(k-2)^j$ and $\mu = (l+1) \cdot k(k-2)^{h-1}$.*

4 Balanced Allocation on Dense Graphs

In this section we show the upper bound for the maximum load attained by $\mathcal{A}(G, l)$ for d-regular graph with $d = \omega(\log n)$. Let us recall the set of parameters for given G and l as follows,

$$\begin{aligned} k &:= \max\{4, \lfloor l/\sqrt{\log_d n} \rfloor\}, \\ \delta &:= \lfloor \lfloor l/k \rfloor / 4 \rfloor, \\ \rho &:= \lceil 8 \log_d n/\delta^2 \rceil, \end{aligned}$$

and $U_{n_1,l,h}$ is the event that at the end of round n_1, there is a nodes with load at least $h\rho+c+1$, where c is a constant and

$$h := \left\lceil \frac{\log\log n}{\log(k-2)} \right\rceil.$$

Note that when $l = (\log n)^{\frac{1+\epsilon}{2}}$ with constant $\epsilon \in (0,1)$, then

$$k = \lfloor l/\sqrt{\log_d n} \rfloor \geqslant l/\sqrt{\log_3 n} \geqslant (\log n)^{\epsilon/3}.$$

Thus, $h = \left\lceil \frac{\log\log n}{\log(k-2)} \right\rceil$ is a constant. Therefore, in order to apply Lemma 3.2 for this case, it is sufficient that G has girth at least $10hl$ or $\omega(l)$. Also we have the following useful lemma whose proof appears in Appendix D.

Lemma 4.1. *With probability $1-o(1/n)$, $\mathcal{N}_\delta$ holds.*

Theorem 4.2. *Suppose that G is a d-regular graph with girth at least $10hl$ and $d = \omega(\log n)$. Then, with high probability the maximum load attained by $\mathcal{A}(G,l)$, denoted by m^*, is bounded from above as follows:*

I. If $\omega(1) \leqslant l \leqslant 4\gamma_G$, where $\gamma_G = \sqrt{\log_d n}$. Then we have

$$m^* = \mathcal{O}\left(\frac{\log_d n \cdot \log\log n}{l^2}\right).$$

II. If $l \geqslant 4\gamma_G$, then we have

$$m^* = \mathcal{O}\left(\frac{\log\log n}{\log(l/\gamma_G)}\right).$$

Note that when $l = \Theta(\gamma_G)$, we get the maximum load $\mathcal{O}(\log\log n)$.

Proof. By Lemma 2.1 we have that $\mathcal{A}(G,l)$ is an (α, n_1)-uniform, where $n_1 = \lfloor n/(6e\alpha) \rfloor$. Let us divide the allocation process into s phases, where s is the smallest integer satisfying $sn_1 \geqslant n$. We now focus on the maximum load attained by $\mathcal{A}$ after allocating n_1 balls in the first phase, which is denoted by m_1^*. Let us assume that $U_{n_1,l,h}$ happens. Now, in order to apply Lemma 3.2, we only need that G has girth at least $10hl$. By Lemma 3.2, if $U_{n_1,l,h}$ conditioning on $\mathcal{N}_\delta$ happens, then there is a c-loaded (λ,μ)-tree T contained in $\mathcal{H}_{n_1}$, where $\lambda = 1 + k\sum_{j=0}^{h-1}(k-2)^j$ and $\mu = (l+1)\cdot k(k-2)^{h-1}$. Thus, we get

$$\begin{aligned}
\mathbf{Pr}\left[U_{n_1,l,h} \mid \mathcal{N}_\delta\right]\mathbf{Pr}\left[\mathcal{N}_\delta\right] &\leqslant \mathbf{Pr}\left[T \text{ exists} \mid \mathcal{N}_\delta\right]\mathbf{Pr}\left[\mathcal{N}_\delta\right] \\
&= \mathbf{Pr}\left[T \text{ exists and } \mathcal{N}_\delta\right] \\
&\leqslant \mathbf{Pr}\left[T \text{ exists}\right].
\end{aligned}$$

Therefore using the law of total probability and the above inequality we have

$$\begin{aligned}
\mathbf{Pr}\left[U_{n_1,l,h}\right] &= \mathbf{Pr}\left[U_{n_1,l,h} \mid \mathcal{N}_\delta\right]\mathbf{Pr}\left[\mathcal{N}_\delta\right] + \mathbf{Pr}\left[U_{n_1,l,h} \mid \neg\mathcal{N}_\delta\right]\mathbf{Pr}\left[\neg\mathcal{N}_\delta\right] \\
&\leqslant \mathbf{Pr}\left[T \text{ exists}\right] + \mathbf{Pr}\left[\neg\mathcal{N}_\delta\right] \\
&= \mathbf{Pr}\left[T \text{ exists}\right] + o(1/n). \qquad (1)
\end{aligned}$$

where the last inequality follows from $\mathbf{Pr}\left[\neg\mathcal{N}_\delta\right] = o(1/n)$ by Lemma 4.1. By definition of h, we get

$$\lambda = 1 + k(1 + (k-2)^h) \leqslant 2k\log n$$

and

$$\mu = (l+1)k(k-2)^{h-1} \geqslant (l+1)(k-2)^h \geqslant (l+1)\log n.$$

It only remains to bound $\mathbf{Pr}\left[T \text{ exists}\right]$. By applying Lemma 2.2 and substituting μ and λ, we conclude that

$$\mathbf{Pr}\left[T \text{ exists}\right] \leqslant n\exp(4\lambda\log(l+1)) - c\mu) \leqslant n\exp\{-z\log n\},$$

where $z = c(l+1) - 8k\log(l+1)$. Depending on k we consider two cases: First, $k = 4$. Then it is easy to see there exists a constant c such that $z \geqslant 2$. Second, $k = \lfloor l/\gamma_G \rfloor$. We know that $l < \log_d n$, so we have $l \leqslant \gamma_G^2$ and hence,

$$z \geqslant cl - 8l\log l/\gamma_G \geqslant l(c - 16\log\gamma_G/\gamma_G) = l(c - o(1)).$$

This yields that for some integer $c > 0$, $z = l(c-o(1)) > 2$ and hence in both cases we get $\mathbf{Pr}\left[T \text{ exists}\right] = o(1/n)$. Now, by Inequality (1) we infer that $m_1^* \leqslant h\rho + c + 1$ with probability $1 - o(1/n)$. In what follows we show the sub-additivity of the algorithm and concludes that

in the second phase the maximum load increases by at most m_1^* whp. Assume that we have a copy of G, say G', whose nodes have load exactly m_1^*. Let us consider the allocation process of a pair of balls (n_1+t,t), for every $0 \leqslant t \leqslant n_1$, by $\mathcal{A}(G,l)$ and $\mathcal{A}(G',l)$. Let $X_u^{n_1+t}$ and Y_u^t, $t \geqslant 0$ denote the load of $u \in V(G) = V(G')$ after allocating balls n_1+t and t by $\mathcal{A}(G,l)$ and $\mathcal{A}(G',l)$, respectively. Now we show that for every integer $0 \leqslant t \leqslant n_1$ and $u \in V(G)$ we have that

$$X_u^{n_1+t} \leqslant Y_u^t. \tag{2}$$

When $t=0$, clearly the inequality holds because $Y_u^0 = m_1^*$. We couple the both allocation processes $\mathcal{A}(G,l)$ and $\mathcal{A}(G',l)$ for a given pair of balls (n_1+t,t), $t \geqslant 0$, as follows. For every $1 \leqslant t \leqslant n_1$, the coupled process first picks a one-to-one labeling function $\sigma_t : V(G) \to \{1,2,\ldots,n\}$ uniformly at random. (Note that σ_t is also defined for G' as $V(G) = V(G')$.) Then it applies $\mathcal{A}(G,l)$ and selects l-walks W_{n_1+t} and its copy, say W'_t, in G'. After that, balls n_1+t and t are allocated on least loaded nodes of W_{n_1+t} and W'_t, respectively, and ties are broken in favor of nodes with minimum label. It is easily checked that the defined process is a coupling. Let us assume that Inequality (2) holds for every $t_0 \leqslant t$, then we show it for $t+1$. Let $v \in W_{n_1+t+1}$ and $v' \in W'_{t+1}$ denote the nodes that are the destinations of pair $(n_1+t+1, t+1)$. Now we consider two cases:

1. $X_v^{n_1+t} < Y_v^t$. Then allocating ball n_1+t+1 on v implies that

   $$X_v^{n_1+t} + 1 = X_v^{n_1+t+1} \leqslant Y_v^t \leqslant Y_v^{t+1}.$$

   So, Inequality (2) holds for $t+1$ and every $u \in V(G)$.

2. $X_v^{n_1+t} = Y_v^t$. Since $W_{n_1+t+1} = W'_{t+1}$, $v \in W'_{t+1}$ and $v' \in W_{n_1+t+1}$. Also we know that v and v' are nodes with minimum load contained in W_{n+t+1} and W_{t+1}, So we have,

   $$X_v^{n_1+t} \leqslant X_{v'}^{n_1+t} \leqslant Y_{v'}^t \leqslant Y_v^t.$$

   Since $Y_v^t = X_v^{n_1+t}$, we have

   $$Y_{v'}^t = Y_v^t = X_v^{n_1+t}.$$

   If $v \neq v'$ and $\sigma_{t+1}(v') < \sigma_{t+1}(v)$, then it contradicts the fact that ball n_1+t+1 is allocated on v. Similarly, if $\sigma_{t+1}(v') > \sigma_{t+1}(v)$, it contradicts that ball t is allocated on v'. So, we have $v = v'$ and

   $$X_v^{n_1+t} + 1 = X_v^{n+t+1} = Y_v^t + 1 = Y_v^{t+1}.$$

So in both cases, Inequality (2) holds for every $t \geqslant 0$. If we set $t = n_1$, then the maximum load attained by $A(G',l)$ is at most $2m_1^*$ whp. Therefore, by Inequality (2), $2m_1^*$ is an upper bound for the maximum load attained by $\mathcal{A}(G,l)$ in the second phase as well. Similarly, we apply the union bound and conclude that after allocating the balls in s phases, the maximum load m^* is at most sm_1^* with probability $1 - o(s/n) = 1 - o(1/n)$. □

Acknowledgment. The author wants to thank Thomas Sauerwald for introducing the problem and several helpful discussions.

References

[1] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9:585–603, 2007.

[2] Noga Alon and Eyal Lubetzky. Poisson approximation for non-backtracking random walks. *Israel J. Math.*, 174(1):227–252, 2009.

[3] Anne Auger and Benjamin Doerr. *Theory of Randomized Search Heuristics: Foundations and Recent Developments*. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2011.

[4] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999.

[5] Petra Berenbrink, André Brinkmann, Tom Friedetzky, and Lars Nagel. Balls into bins with related random choices. *J. Parallel Distrib. Comput.*, 72(2):246–253, 2012.

[6] Paul Bogdan, Thomas Sauerwald, Alexandre Stauffer, and He Sun. Balls into bins via local search. In *Proc. 24th Symp. Discrete Algorithms (SODA)*, pages 16–34, 2013.

[7] John W. Byers, Jeffrey Considine, and Michael Mitzenmacher. Geometric generalizations of the power of two choices. In *Proc. 16th Symp. Parallelism in Algorithms and Architectures (SPAA)*, pages 54–63, 2004.

[8] Colin Cooper, Alan M. Frieze, and Tomasz Radzik. Multiple random walks in random regular graphs. *SIAM J. Discrete Math.*, 23(4):1738–1761, 2009.

[9] Xavier Dahan. Regular graphs of large girth and arbitrary degree. *Combinatorica*, 34(4):407–426, 2014.

[10] Brighten Godfrey. Balls and bins with structure: balanced allocations on hypergraphs. In *Proc. 19th Symp. Discrete Algorithms (SODA)*, pages 511–517, 2008.

[11] Krishnaram Kenthapadi and Rina Panigrahy. Balanced allocation on graphs. In *Proc. 17th Symp. Discrete Algorithms (SODA)*, pages 434–443, 2006.

[12] Donald Knuth. *The Art of Computer Programming, Vol. 1: Fundamental Algorithms*. Adison-Wesley, third edition, 1997.

[13] Micheal Mitzenmacher, Andréa W. Richa, and Ramesh Sitaraman. The power of two random choices: A survey of technique and results. *In Handbook of Randomized Computation Volume 1*, pages 255–312, 2001.

[14] Yuval Peres, Kunal Talwar, and Udi Wieder. Graphical balanced allocations and the $(1+\beta)$-choice process. *Random Struct. Algorithms*, DOI: 10.1002/rsa.20558, 2014.

[15] Berthold Vöcking. How asymmetry helps load balancing. *J. ACM*, 50(4):568–589, 2003.

A Omitted Proofs of Section 2

In this section we show some useful results about interference graph $\mathcal{I}(G,l)$ and present the omitted proofs of Section 2.

Lemma A.1. *Suppose that $V(\mathcal{I})$ and $\Delta(\mathcal{I})$ denote the vertex set and the maximum degree of interference graph $\mathcal{I}(G,l)$, respectively. Then we have,*

(i) $|V(\mathcal{I})| = nd(d-1)^{l-1}/2$,

(ii) $\Delta(\mathcal{I}) \leqslant (l+1)^2 d(d-1)^{l-1}$.

Furthermore, the number of rooted λ-vertex trees contained in $\mathcal{I}$ is bounded by $4^\lambda \cdot |V(\mathcal{I})| \cdot \Delta(\mathcal{I})^{\lambda-1}$.

Proof. It is easy to see that in a graph with girth at least $\omega(l)$, the number of l-walks is exactly $nd(d-1)^{l-1}/2$, (without ordering) which is the size of $V(\mathcal{I})$. Since the graph locally looks like a d-ary tree, the total number of l-walks including v as j-th visited node is at most

$$d(d-1)^{j-2}(d-1)^{l-j-1} = d(d-1)^{l-1}.$$

Index j varies from 0 to l, so v can be an element of at most $(l+1)d(d-1)^{l-1}$ l-walks. Also, every l-walk contains $l+1$ elements and hence every l-walk intersects at most $(l+1)^2 d(d-1)^{l-1}$ other l-walks. Thus we get

$$\Delta(\mathcal{I}) \leqslant (l+1)^2 \cdot d(d-1)^{l-1}.$$

Let us now bound the total number of rooted λ-vertex trees contained in $\mathcal{I}$. It was shown that the total number of different shape rooted trees on λ vertices is 4^λ (For example see [12]); we say two rooted trees have different shapes if they are not isomorphic. For any given shape, there are $|V(\mathcal{I})|$ ways to choose the root. As soon as the root is fixed, each vertex in the first level can be chosen in at most $\Delta(\mathcal{I})$ many ways. By selecting the vertices of the tree level by level we have that each vertex except the root can be chosen in at most $\Delta(\mathcal{I})$ ways. So the total number of rooted λ-vertex trees in $\mathcal{I}$ is bounded by

$$4^\lambda \cdot |V(\mathcal{I})| \cdot \Delta(\mathcal{I})^{\lambda-1}.$$

□

Corollary A.2. *The size of family of (λ,μ)-trees is bounded by $4^\lambda |V(\mathcal{I})| \Delta(I)^{\lambda-1}$.*

Proof. We know that every (λ,μ)-tree T is a rooted λ-vertex subtree of $\mathcal{I}$ with the additional property that $|\cup_{W\in V(T)} W| \geqslant \mu$. This implies that the size of family of rooted λ-vertex subtrees of $\mathcal{I}$ is an upper bound for the size of family of (λ,μ)-trees and hence by applying Lemma A.1, we reach the upper bound $4^\lambda |V(\mathcal{I})| \Delta(I)^{\lambda-1}$. □

A.1 Proof of the Key Lemma

In this subsection we first present several useful lemmas and then show the key lemma 2.1. Before that let us define some notations. For every $S \subseteq V(G)$, $\texttt{Empty}_t(S)$ denotes the number of empty nodes contained in S after allocating t balls. Let $N(v)$ denote the set of neighbors of v. Note that to avoid a lengthy case analysis we do not optimize the constants.

Lemma A.3 (Deviation bounds for moderate independency)**.** *Let $X_1, \cdots, X_n$ be arbitrary binary random variables. Let $X_1^*, X_2^*, \cdots, X_n^*$ be binary random variables that are mutually independent and such that for all i, X_i^*, is independent of $X_1, \cdots, X_{i-1}$. Assume that for all i and all $x_1, ..., x_{i-1} \in \{0, 1\}$,*

$$\mathbf{Pr}\left[X_i = 1 | X_1 = x_1, \cdots, X_{i-1} = x_{i-1}\right] \geqslant \mathbf{Pr}\left[X_i^* = 1\right].$$

Then for all $a \geqslant 0$, we have

$$\mathbf{Pr}\left[\sum_{i=1}^{n} X_i \leqslant a\right] \leqslant \mathbf{Pr}\left[\sum_{i=1}^{n} X_i^* \leqslant a\right]$$

and the latter term can be bounded by any deviation bound for independent random variables.

The proof of the above lemma can be found in [3, Lemma 1.18].

Lemma A.4. *Suppose that $\mathcal{A}(G, l)$ has allocated the balls until the $(t+1)$-th ball, for some $0 \leqslant t \leqslant n$. Then, for every given $v \in V(G)$ we have*

$$\mathbf{Pr}\left[v \in W_{t+1}\right] = \sum_{i=0}^{l} \mathbf{Pr}\left[C_i\right] = (l+1)/n,$$

where C_i, $0 \leqslant i \leqslant l$, is the event that for $W_{t+1} = (u_0, u_1, \ldots, u_l)$, we have $v = u_i$. Furthermore, for every $0 \leqslant i \leqslant l$, we have

$$\mathbf{Pr}\left[C_i\right] = 1/n.$$

Proof. Let us fix an arbitrary $0 \leqslant i \leqslant l$ and any $v \in V(G)$. Since G has girth at least $\omega(l)$ and locally looks like a d-regular tree, we can easily compute the number of l-walks visiting v in the i-th step, that is

$$d(d-1)^{i-2} \times (d-1)^{l-(i-1)} = d(d-1)^{l-1}.$$

On the other hand in each round, $\mathcal{A}(G, l)$ picks an l-walk randomly from $nd(d-1)^{l-1}$ possible l-walks. Thus, we get

$$\mathbf{Pr}\left[C_i\right] = \frac{d(d-1)^{l-1}}{nd(d-1)^{l-1}} = \frac{1}{n}$$

and

$$\sum_{i=0}^{l} \mathbf{Pr}\left[C_i\right] = \sum_{i=0}^{l} \frac{1}{n} = \frac{l+1}{n}.$$

□

Lemma A.5. *Suppose that with probability $1 - o(n^{-2})$, for every $u \in V(G)$,*

$$\texttt{Empty}_t(N(u)) \geqslant |N(u)|/2 = d/2.$$

Then for every $v \in V(G)$,

$$\mathbf{Pr}\left[\textit{ball } t+1 \textit{ is allocated on } v \textit{ by } \mathcal{A}\ \right] \leqslant \frac{\alpha}{n},$$

where α is a constant.

Proof. Let $\mathcal{E}_{t+1,v}$ be the event that ball $t+1$ is placed on a given node $v \in V(G)$ and $\mathcal{F}_{t+1}$ be the event that at least $l/10$ of nodes in W_{t+1} are empty. Let $\neg\mathcal{F}_{t+1}$ denote the negation of $\mathcal{F}_{t+1}$. Using the law of total probability, for every $v \in V(G)$ we have

$$\begin{aligned}
\mathbf{Pr}\left[\mathcal{E}_{t+1,v}\right] &= \underbrace{\mathbf{Pr}\left[\mathcal{E}_{t+1,v} | v \notin W_{t+1}\right] \cdot \mathbf{Pr}\left[v \notin W_{t+1}\right]}_{=0} \\
&\quad + \underbrace{\mathbf{Pr}\left[\mathcal{E}_{t+1,v} | v \in W_{t+1} \text{ and } \mathcal{F}_{t+1}\right] \cdot \mathbf{Pr}\left[v \in W_{t+1} \text{ and } \mathcal{F}_{t+1}\right]}_{\leqslant (10/l)\,\mathbf{Pr}\left[v \in W_{t+1}\right]} \\
&\quad + \underbrace{\mathbf{Pr}\left[\mathcal{E}_{t+1,u} | v \in W_{t+1} \text{ and } \neg\mathcal{F}_{t+1}\right]}_{\leqslant 1} \cdot \mathbf{Pr}\left[v \in W_{t+1} \text{ and } \neg\mathcal{F}_{t+1}\right] \\
&\leqslant \frac{10}{l} \cdot \mathbf{Pr}\left[v \in W_{t+1}\right] + \mathbf{Pr}\left[v \in W_{t+1} \text{ and } \neg\mathcal{F}_{t+1}\right]
\end{aligned}$$

where the first summand follows since if $v \notin W_{t+1}$, then ball $t+1$ cannot be placed on v and the second one follows because ties are broken uniformly at random. Now, by applying Lemma A.4 and and Bayse' rule we have,

$$\begin{aligned}\mathbf{Pr}\left[\mathcal{E}_{t+1,v}\right] &\leqslant \frac{10(l+1)}{ln} + \sum_{i=0}^{l} \mathbf{Pr}\left[\ \neg\mathcal{F}_{t+1} \text{and } C_i\right] \\ &= \frac{10(l+1)}{ln} + \sum_{i=0}^{l} \mathbf{Pr}\left[\ \neg\mathcal{F}_{t+1} \mid C_i\right](1/n). \end{aligned} \tag{3}$$

In what follows we will show that for every i,

$$\mathbf{Pr}\left[\neg\mathcal{F}_{t+1} | C_i\right] \leqslant 12/l.$$

Plugging the above bound in Inequity (3) yields that for every $v \in V(G)$,

$$\mathbf{Pr}\left[\mathcal{E}_{t+1,v}\right] \leqslant \frac{22(l+1)}{ln},$$

where $22(l+1)/l$ is indeed a constant and hence the statement is proved.

Conditioning on event C_i, we only know that node v is the i-th visited node in W_{t+1}, for some $0 \leqslant i \leqslant l$. Clearly, W_{t+1} can be viewed as the union of two edge-disjoint NBRWs of lengths $(i-1)$ and $l-(i-1)$ started from v, namely W_v^1 and W_v^2. Without loss of generality, assume that $|V(W_v^1)| = s \geqslant 2$ and let

$$W_v^1 = (v = u_1, u_2, \ldots, u_s),$$

where $d(v,u_j) < d(v,u_{j'})$ for every $1 < j < j' \leqslant s$. Clearly, every $u_j \in W_v^1$, $2 \leqslant j \leqslant s$, is randomly chosen from a subset of $N(u_{j-1})$, say $S_j \subseteq N(u_{j-1})$ (because we run a NBRW from u_{j-1} to reach u_j). If it happens that the NBRW has already traversed edge $\{w, u_{j-1}\}$, for some node w, then the walk cannot take this edge again and hence

$$|S_j| \in \{d, d-1\}.$$

Let us define an indicator random variable X_{u_j} for every u_j, $2 \leqslant j \leqslant s$, which takes one whenever u_j is empty and zero otherwise. Thus we have

$$\mathbf{Pr}\left[X_{u_j} = 1\right] = \frac{\texttt{Empty}(S_j)}{|S_j|}.$$

Let $\mathcal{K}_j$, $2 \leqslant j \leqslant s$, denote the event that the number of empty nodes of $N(u_{j-1})$ is at least $d/2$. By the assumption, we have $\mathbf{Pr}\left[\mathcal{K}_j\right] = 1 - o(n^{-2})$. So, for every u_j, $2 \leqslant j \leqslant s$, we get

$$\begin{aligned}\mathbf{Pr}\left[X_{u_j} = 1\right] &= \mathbf{Pr}\left[X_{u_j} = 1 \mid \mathcal{K}_j\right]\mathbf{Pr}\left[\mathcal{K}_j\right] + \mathbf{Pr}\left[X_{u_j} = 1 \mid \neg\mathcal{K}_j\right]\mathbf{Pr}\left[\neg\mathcal{K}_j\right] \\ &\geqslant 1/2((d-2)/(d-1))(1-o(n^{-2})) + o(n^{-2}) \geqslant 1/3,\end{aligned}$$

where the first inequality follows from

$$\frac{\texttt{Empty}(S_j)}{|S_j|} \geqslant \min\left\{\frac{d/2}{d}, \frac{d/2-1}{d-1}\right\} = \frac{d-2}{2(d-1)}$$

Since the above lower bound is independent of any X_{u_j}, $2 \leqslant j' \leqslant j$, we have that for every $2 \leqslant j \leqslant s$,

$$\mathbf{Pr}\left[X_{u_j} = 1 \mid X_{u_1} = x_1, \cdots, X_{u_{j-1}} = x_{j-1}\right] \geqslant 1/3.$$

A similar argument also works for the nodes visited by W_{t+1}^2 and we get $\mathbf{Pr}\left[X_u = 1\right] \geqslant 1/3$, for every $u \in W_{t+1} \setminus \{v\}$. Let $Y = \sum_{u \in W_{t+1}\setminus\{v\}} X_u$ be the number of empty nodes in $W_{t+1} \setminus \{v\}$. Then, we have that $\mathbf{E}\left[Y\right] \geqslant l/3$. Let Y^* be the summation of l independent Bernoulli random variables with success probability $1/3$. By applying Lemma A.3 we get,

$$\begin{aligned}\mathbf{Pr}\left[\neg\mathcal{F}_{t+1} | C_i\right] &\leqslant \mathbf{Pr}\left[Y < l/6\right] \\ &\leqslant \mathbf{Pr}\left[Y^* < \mathbf{E}\left[Y^*\right]/2\right] \leqslant \mathbf{Pr}\left[|Y^* - \mathbf{E}\left[Y^*\right]| \geqslant \mathbf{E}\left[Y^*\right]/2\right].\end{aligned}$$

We know that $\mathbf{Var}\left[Y^*\right] \leqslant \mathbf{E}\left[Y^*\right]$, so applying Chebychev's bound results into

$$\mathbf{Pr}\left[|Y^* - \mathbf{E}\left[Y^*\right]| \geqslant \mathbf{E}\left[Y^*\right]/2\right] \leqslant \frac{\mathbf{Var}\left[Y^*\right]}{(\mathbf{E}\left[Y^*\right]/2)^2} \leqslant \frac{4}{\mathbf{E}\left[Y^*\right]}.$$

Thus, we get

$$\mathbf{Pr}\left[\neg\mathcal{F}_{t+1} | C_i\right] \leqslant 4/\mathbf{E}\left[Y^*\right] \leqslant 12/l.$$

□

In order to prove our key lemma, we apply a potential function argument which is similar to [6, Theorem 1.4].

Proof of Key Lemma 2.1. Let us define potential function

$$\Phi(t) = \sum_{u \in V(G)} \exp(a_t(u)),$$

where $a_t(u)$ denotes the number of nonempty nodes of $N(u)$ after allocating t balls. It is clear that $\Phi(0) = n$. Let us assume that after allocating t balls we have $\Phi(t) \leqslant n \cdot e^{d/4}$.

$$\mathrm{e}^{a_t(u)} \leqslant \Phi(t) \leqslant \mathrm{e}^{\log n + d/4}.$$

Since $d = \omega(\log n)$, we get

$$a_t(u) \leqslant \log n + d/4 < d/2$$

and consequently $\texttt{Empty}_t(N(u)) \geqslant \frac{d}{2}$, for every $u \in V(G)$. Let us define indicator random variable $I_{t+1}(u)$ for every $u \in V(G)$ as follows:

$$I_{t+1}(u) := \begin{cases} 1 & \text{if ball } t+1 \text{ is placed on an empty node in } N(u), \\ 0 & \text{otherwise.} \end{cases}$$

Applying Lemma A.5 shows that if $\texttt{Empty}_t(N(u)) \geqslant \frac{d}{2}$, then for every $u \in V(G)$,

$$\mathbf{Pr}\left[I_{t+1}(u) = 1\right] \leqslant \frac{\alpha \cdot \texttt{Empty}_t(N(u))}{n} \leqslant \frac{\alpha \cdot d}{n},$$

where α is a constant. So we get

$$\begin{aligned} &\mathbf{E}\left[\Phi(t+1) \mid \Phi(t) \leqslant n \cdot e^{d/4}\right] \\ &\leqslant \sum_{u \in V(G)} \left\{\mathbf{Pr}\left[I_{t+1}(u) = 1\right] \cdot e^{a_t(u)+1} + \mathbf{Pr}\left[I_{t+1}(u) = 0\right] \cdot e^{a_t(u)}\right\} \\ &\leqslant \sum_{u \in V(G)} \left(1 + \frac{\alpha \cdot \mathrm{e} \cdot d}{n}\right) \cdot e^{a_t(u)} = \left(1 + \frac{\alpha \cdot \mathrm{e} \cdot d}{n}\right) \Phi(t). \end{aligned}$$

Let us define $\Psi(t) := \min\{\Phi(t), n \cdot e^{\Delta/4}\}$. By using above recursive inequality we have that

$$\mathbf{E}\left[\Psi(t+1)\right] \leqslant \left(1 + \frac{\alpha \cdot \mathrm{e} \cdot d}{n}\right) \Psi(t).$$

Thus, inductively we have that $\mathbf{E}\left[\Psi(t)\right] \leqslant \left(1 + \frac{\alpha \cdot \mathrm{e} \cdot d}{n}\right)^t \Psi(0)$. Let us define $n_1 = n/(6\mathrm{e}\alpha)$. Then applying Markov's inequality implies that

$$\mathbf{Pr}\left[\Psi(n_1) \geqslant n \cdot e^{d/4}\right] \leqslant \frac{\left(1 + \frac{\alpha \cdot \mathrm{e} \cdot d}{n}\right)^{n_1}}{e^{d/4}} \leqslant e^{-d/12}$$

So with probability $1 - n^{-\omega(1)}$, we have $\Phi(n_1) = \Psi(n_1) < n \cdot e^{d/4}$. Since $\Phi(t)$ is an increasing function in t, we have that $\Phi(t) \leqslant n \cdot e^{d/4}$, for every $0 \leqslant t \leqslant n_1$, and hence with probability $1 - o(n^{-2})$, for every $u \in V(G)$,

$$\texttt{Empty}_t(N(u)) \geqslant d/2.$$

So, applying Lemma A.5 shows that for every $0 \leqslant t \leqslant n_1$ and $u \in V(G)$,

$$\mathbf{Pr}\left[\text{ball } t+1 \text{ is placed on } u \text{ by } \mathcal{A}(G, l)\right] \leqslant \frac{\alpha}{n}.$$

□

A.2 Proof of Lemma 2.2

Proof. Let us fix an arbitrary (λ,μ)-tree $T \subseteq \mathcal{I}(G,l)$ and p_1 be the probability that using λ balls T is built and contained in $\mathcal{H}_{n_1}$. There are at most $n_1 \leqslant n$ ways to choose one ball per vertex of T and hence at most n^λ ways to choose λ balls that are going to pick the vertices of T. On the other hand, every ball picks a given vertex of T with probability $1/|V(\mathcal{I})|$. Thus we get,

$$p_1 \leqslant n^\lambda \cdot (1/V(\mathcal{I}))^\lambda.$$

Now, we have to add c additional balls for very node in $\cup_{W\in V(T)}W$, where $|\cup_{W\in V(T)} W| = \mu+q$, for some integer $q \geqslant 0$. Let p_2 denote the probability that such a event happens. Since $\mathcal{A}(G,l)$ is (α,n_1)-uniform with $n_1 = \lfloor n/(6\mathrm{e}\alpha)\rfloor$, we get

$$\begin{aligned} p_2 &\leqslant \sum_{q=0}^{\infty} \binom{n_1}{c\cdot(\mu+q)} \left(\frac{\alpha\cdot(\mu+q)}{n}\right)^{c\cdot(\mu+q)} \\ &\leqslant \sum_{q=0}^{\infty} \left(\frac{\mathrm{e}\cdot n_1}{c\cdot(\mu+q)}\right)^{c\cdot(\mu+q)} \cdot \left(\frac{\alpha\cdot(\mu+q)}{n}\right)^{c\cdot(\mu+q)} \\ &\leqslant \sum_{q=0}^{\infty} \left(\frac{n_1\cdot\alpha\cdot\mathrm{e}}{n\cdot c}\right)^{c\cdot(\mu+q)} = (1/6c)^{c\mu}\sum_{q=0}^{\infty}(1/6c)^{cq} \\ &\leqslant 2\cdot(1/6c)^{c\mu}, \end{aligned}$$

where we use the fact that for integers $1 \leqslant a \leqslant b$, $\binom{b}{a} \leqslant (\frac{\mathrm{e}b}{a})^a$ and the last inequality follows from $\sum_{q=0}^{\infty}(1/6c)^{cq} \leqslant 2$. Since balls are mutually independent, $p_1\cdot p_2$ is an upper bound for the probability that c-loaded (λ,μ)-tree T appears in $\mathcal{H}_{n_1}$. By Corollary A.2 we have an upper bound for the size of family of all (λ,μ)-trees. Hence, taking the union bound over all (λ,μ)-trees gives an upper bound for appearance probability of a c-loaded (λ,μ)-tree in $\mathcal{H}_{n_1}$. Thus we get,

$$\begin{aligned} 4^\lambda|V(\mathcal{I})|\cdot\Delta^{\lambda-1}\cdot p_1\cdot p_2 &\leqslant 2\cdot 4^\lambda|V(\mathcal{I})|\cdot\Delta^{\lambda-1}\left(\frac{n}{V(\mathcal{I})}\right)^\lambda\cdot\left(\frac{1}{6c}\right)^{c\cdot\mu} \\ &\leqslant 2n\cdot 4^\lambda\cdot\left(\frac{\Delta(\mathcal{I})}{V(\mathcal{I})}\right)^{\lambda-1}\cdot\left(\frac{1}{6\cdot c}\right)^{c\cdot\mu}. \end{aligned}$$

By Lemma A.1 we have $|V(\mathcal{I})| = nd(d-1)^{l-1}/2$, $\Delta(\mathcal{I}) \leqslant (l+1)^2 d(d-1)^{l-1}$. So the above bound is simplified as follows,

$$2n\cdot 4^\lambda\left(2(l+1)^2\right)^{\lambda-1}\left(\frac{1}{6}\right)^{c\cdot\mu} \leqslant n(l+1)^{4\lambda}6^{-c\mu} \leqslant n\exp(4\lambda\log(l+1) - c\mu),$$

where the first inequality follows from

$$2\cdot 4^\lambda\cdot 2^{\lambda-1}(l+1)^{2(\lambda-1)} = 8^\lambda(l+1)^{2(\lambda-1)} \leqslant (l+1)^{4\lambda},$$

which is true for every $l \geqslant 2$. □

B Proof of Lemma 3.1

Proof. Let us fix an arbitrary subpath $P_i = [u_i,u_{i+1}] \in \mathcal{P}_k(W)$ and partition it into 3 edge-disjoint subpaths, say $P_i = [u_i,u] \cup_g [u,v] \cup_g [v,u_{i+1}]$, such that

$$d(u_i,u) = d(v,u_{i+1}) = \delta,$$

where $\delta = \lfloor\lfloor l/k\rfloor/4\rfloor$. By the Partition step in $PB(\rho,k)$, we know that

$$d(u_i,u_{i+1}) \in \{\lfloor l/k\rfloor, \lceil l/k\rceil\}.$$

So we have

$$d(u,v) = d(u_i,u_{i+1}) - 2\delta \geqslant 4\delta - 2\delta = 2\delta.$$

Let $S = W \cap V([u,v])$ and $B(S)$ denotes the set of all balls allocated on nodes of S at height at least $f(W)-\rho$. Let W_t denote the chosen l-walk by ball $t \in B(S)$. Since each ball $t \in B(S)$ was allocated on a node of S at height at least $f(W)-\rho$, we have that W_t intersects $[u,v] \subset P_i$,

$f(W_t) \geqslant f(W) - \rho$ and $W_t \cap W \neq \emptyset$. So each W_t satisfies (C2). Now, among all W_t, $t \in B(S)$, we find an l-walk that satisfies (C1) as well. We have

$$|S| \geqslant 2\delta,$$

Every node in S has load at least $f(W) \geqslant \rho + 1$. So, every node in S has at least ρ balls at height at least $f(W) - \rho \geqslant 1$. Therefor we have,

$$\begin{aligned} |B(S)| \geqslant |S|\rho \geqslant (2\delta)\rho &\geqslant (2\delta)(6\log_{d-1} n/\delta^2) \\ &= 12\log_{d-1} n/\delta. \end{aligned}$$

By using the above inequality we have,

$$|\{W_t, t \in B(S)\}| = |B(S)| \geqslant 12\log_{d-1} n/\delta.$$

Recall that $P_i = [u_i, u] \cup [u, v] \cup [v, u_{i+1}]$. If for some $t \in B(S)$, W_t contains u_i (or u_{i+1}), then it also contains subpath $[u_i, u]$ (or $[v, u_{i+1}]$), because W_t intersects $[u, v]$ and G has girth $\omega(l)$. Conditioning on $\mathcal{N}_\delta$, $[u_i, u]$ and $[v, u_{i+1}]$ are contained in less than $12\log_{d-1} n/\delta$ l-walks. So the above inequality shows that there is at least one ball, say $t_0 \in B(S)$, whose corresponding l-walk W_{t_0} contains neither u_i nor u_{i+1} and thus it satisfies $(C1)$. Therefore we conclude that, for each $P_i \in \mathcal{P}_k(W)$, W_{P_i} exists and $PB(\rho, k)$ on W is valid. □

C Proof of Lemma 3.2

Before we present the proof of Lemma 3.2, we need to show some lemmas about the properties of the recursive construction of the witness tree. Suppose that $H_j \subset G$, $0 \leqslant j \leqslant h$, be the graphical union of all l-walks up to $j+1$-the level (i.e., $\mathcal{L}_{j+1}$). Then we have the following lemma.

Lemma C.1. *If G has girth at least $10hl$, then, for every $0 \leqslant j \leqslant h$, H_j is a tree.*

Proof. When $j = 0$, clearly $H_0 = R$, where R is the root. So the diameter of H_0 is l. Assume that for some j_0, $0 \leqslant j_0 < h$, the diameter of H_{j_0} is at most $(2j_0 + 1)l$. We know that every l-walk in the $(j_0 + 1)$-th level intersects a path in H_{j_0} so the distance between any two nodes of H_{j_0+1} increases by at most $2l$ and thus the diameter of H_{j_0+1} is at most

$$(2j_0 + 1)l + 2l = (2(j_0 + 1) + 1)l.$$

So we inductively conclude that, for every $0 \leqslant j \leqslant h$, H_j has diameter at most $(2j + 1)l$. If for some j, $0 \leqslant j \leqslant h$, H_j contains a cycle, then the length of the cycle is at most $2 \cdot \texttt{diam}(H_i) \leqslant 2(2j + 1)l \leqslant 6hl$ which contradicts the fact that $H_j \subset G$ and G has girth at least $10hl$. □

Lemma C.2. *For every $1 \leqslant j \leqslant h$, the j-th level contains $k(k-2)^{j-1}$ disjoint l-walks. Moreover every l-walk in the j-the level only intersects one l-walk in the previous levels, which is its father.*

Proof. Let us begin with $j = 1$. For the sake of a contradiction, assume that $W_{P_i}, W_{P_{i'}} \in \mathcal{L}_1$ intersect each other. Recall that the l-walks are created by the Branch step over the edge-disjoint paths, say $P_i = [u_{i_1}, u_{i+1}]$ and $P_{i'} = [u_{i'}, u_{i'+1}] \in \mathcal{P}_k(R)$. Clearly, $W_{P_i} \cup_g W_{P_{i'}}$ is a connected graph as they intersect each other. So, by Condition (C1), in the Partition-Branch procedure, we choose two arbitrary nodes $z \in V(P_i) \cap W_{P_i}$ and $z' \in V(P_{i'}) \cap W_{P_{i'}}$. Also, let $\{u_i, u_{i+1}\}$ and $\{u_{i'}, u_{i'+1}\}$ be the *boundary* of P_i and $P_{i'}$, respectively. Since H_0 is a tree, there is a unique path, say $Q_{z,z'}$, in $H_0 = R$ connecting z to z'. Nodes z and z' have degree 2 in H_0, so $Q_{z,z'}$ contains nodes from boundaries of P_i and $P_{i'}$. By (C1), W_{P_i} and $W_{P_{i'}}$ excludes the boundaries. Thus we get a path from z to z' via $W_{P_i} \cup_g W_{P_{i'}} \subset H_1$ that excludes the boundaries. This contradicts the fact that there is a unique path in $H_1 \supset H_0$, because H_1 is a tree by Lemma C.1. So we infer that there are k disjoint l-walks in $\mathcal{L}_1$ and they only intersect their father (i.e., R). Recall that $P \in \mathcal{P}_k(W)$ is a *free* subpath if it does not share any node with W's father. Since W's, $W \in \mathcal{L}_1$, are mutually disjoint, the nodes contained in the set of free subpaths, $\mathcal{P}_k^0(W)$, for each $W \in \mathcal{L}_1$, have degree at most 2 in H_1, which we call it $\mathcal{D}_1$ property. In other word, $\mathcal{D}_1$ property says that any path in H_1 between nodes of two free subpaths in the first level includes nodes from boundaries of the subpaths (see Figure 2).

Suppose that for some j_0, $1 \leqslant j_0 \leqslant h$, the statement of the lemma and $\mathcal{D}_{j_0}$ hold. Then we show them for the next level as well.

Similar to case $j = 1$, toward a contradiction assume that two l-walks $W_P, W_{P'} \in \mathcal{L}_{j_0+1}$ intersect each other. Then, by $(C2)$ we get a path in $W_P \cup_g W_{P'} \subset H_{j_0+1}$ excluding the boundaries of P and P' that connects a node, say x, from P to another node, say y, in P'. By $\mathcal{D}_{j_0}$ property, the path in H_{j_0} connecting x to y uses nodes from the boundaries, while we get a path in H_{j_0+1} that exclude boundaries. This is a contradiction because $H_{j_0+1} \supset H_{j_0}$ is a tree by Lemma C.1. So the l-walls in $\mathcal{L}_{j_0+1}$ are disjoint and by the construction we have $|\mathcal{L}_{j_0+1}| = (k-2)|\mathcal{L}_{j_0}|$ and hence $|\mathcal{L}_{j_0+1}| = k(k-2)^{j_0}$. It only remains to prove every l-walk only intersect its father in previous levels. Toward a contradiction assume that $W_P \in \mathcal{L}_{j_0+1}$ intersects a path, say W, in previous levels that is not its father. Let $z' \in W_P \cap W$ and $z \in W_P \cap V(P) \subset V(P)$ where $P = [u, v] \in \mathcal{P}_k^0(W')$ and W' is the father of W_P. By $(C2)$, z is neither u nor v. We now get a new path from z to z' in H_{j_0+1} excluding u and v (via $W_p \cup_g W$) that contradicts the fact that there is only one path from z to z' in H_{j_0} including a node from the boundary of P, as $\mathcal{D}_{j_0}$ property holds. We showed that every two l-walks in $\mathcal{L}_{j_0+1}$ are disjoint, so the $\mathcal{D}_{j_0+1}$ holds as well. □

Proof of Lemma 3.2. Let us consider a graph T whose nodes are the set of all l-walks in $\bigcup_{j=0}^{h} \mathcal{L}_j$, where $\mathcal{L}_0 = \{R\}$. And two nodes are connected, if and only if the corresponding l-walks intersect each other or vice versa. By Lemma C.2 for every $1 \leqslant j \leqslant h$, the j-th, level contains $k(k-2)^{j-1}$ disjoint l-walks and they intersect either their fathers or their $k-2$ children. This implies that T is a subtree of interference graph $\mathcal{H}_{n_1}$ with

$$|V(T)| = \lambda = 1 + k\sum_{j=0}^{h-1}(k-2)^j.$$

If we only consider the h-th level, then we get

$$\left|\cup_{W \in V(T)} W\right| \geqslant \mu = (l+1) \cdot k(k-2)^{h-1}.$$

By (C2) in the Partition-Branch procedure we have that for every $W \in \mathcal{L}_j$, $1 \leqslant j \leqslant h$,

$$f(W) \geqslant (h-j)\rho + c.$$

Hence every node in $\cup_{W \in V(T)} W$ has load at least c. □

D Proof of Lemma 4.1

Proof. Let us fix an arbitrary path $[u, v]$ of length δ, where

$$\delta = \lfloor \lfloor l/k \rfloor / 4 \rfloor = \min\{\lfloor l/16 \rfloor, \lfloor \sqrt{\log_d n}/4 \rfloor\}.$$

The latter equality is true because we set $k = \max\{4, \lfloor l/\sqrt{\log_d n} \rfloor\}$. Clearly, if W be an l-walk and $[u, v] \subseteq W = [u_0, u_l]$, then

$$d(u_0, u) + d(v, u_l) = l - \delta.$$

Moreover, G is a d-regular graph with girth at least $\omega(l)$, so the total number of different paths of length l containing $[u, v]$ is

$$\sum_{a+b=l-\delta} (d-1)^a (d-1)^b = (l - \delta + 1) \cdot (d-1)^{l-\delta}.$$

On the other hand the total number of different paths of length l is $n \cdot d \cdot (d-1)^{l-1}/2$. So the probability that in some round t, $1 \leqslant t \leqslant n$, we get $[u, v] \subseteq W_t$ is at most

$$\frac{2(l-\delta+1)(d-1)^{l-\delta}}{n \cdot d \cdot (d-1)^{l-1}} = \frac{2(l-\delta+1)(d-1)}{n \cdot d \cdot (d-1)^{\delta}} \leqslant \frac{2l}{n(d-1)^{\delta}}.$$

Let $u_\delta = \lceil 6 \log_{d-1} n/\delta \rceil$ and $\{t_1, t_2, \ldots, t_{u_\delta}\} \subset [n]$ be a sequence of distinct rounds of size u_δ. We define indicator random variable $X_{t_1, t_2, \ldots, t_{u_\delta}}([u, v])$, which takes one if $[u, v] \subseteq W_{t_i}$, for every $1 \leqslant i \leqslant u_\delta$, and zero otherwise. Thus we get

$$\begin{aligned}
\mathbf{Pr}\left[X_{t_1, t_2 \ldots, t_{u_\delta}}([u, v]) = 1\right] &\leqslant \left(2l/n(d-1)^{\delta}\right)^{u_\delta} \\
&= n^{-u_\delta} \, (d-1)^{(\log_{d-1}(2l) - \delta) u_\delta} \\
&\leqslant n^{-u_\delta} (d-1)^{-u_\delta \cdot \delta/2} = n^{-u_\delta} n^{-3},
\end{aligned}$$

where the last inequality follows from $l \in [\omega(1), o(\log n)]$ and hence,

$$\log_{d-1}(2l) \leqslant \delta/2.$$

There are at most n^{u_δ} sequences of rounds of size u_δ and at most $n(d-1)^{\delta-1}$ paths of length δ. Thus, by using the previous upper bound and the union bound over all sequences of rounds and paths of length δ we have

$$\begin{aligned}
&\sum_{\delta\text{-path}} \sum_{t_1,t_2,\ldots,t_{u_\delta}} \mathbf{Pr}\left[X_{t_1,t_2\ldots,t_{u_\delta}}([u,v]) = 1\right] \\
&\leqslant nd(d-1)^{\delta-1} n^{u_\delta}\, \mathbf{Pr}\left[X_{t_1,t_2\ldots,t_{u_\delta}}([u,v]) = 1\right] \\
&\leqslant o(n^2) n^{u_\delta}\, \mathbf{Pr}\left[X_{t_1,t_2\ldots,t_{u_\delta}}([u,v]) = 1\right] = o(1/n),
\end{aligned}$$

where the last inequality follows from $\delta \leqslant l = o(\log_d n)$. This implies that with probability $1 - o(1/n)$ there is no path of length δ contained in at least u_δ l-walks or equivalently $\mathcal{N}_\delta$ holds. □

E Balanced Allocation on Sparse Graphs

In this section we present allocation algorithm $\mathcal{A}'(G,l)$ for d-regular graphs, with $d \in [3, \mathcal{O}(\log n)]$. The algorithm proceeds as follows: In each round, every ball picks a node uniformly at random and it takes a NBRW of length $l \cdot r_G$ from the chosen node, where $r_G = \lceil 2\log_{d-1}\log n \rceil$. After that the ball collects the load information every r_G-th visited node, called the *potential choice*, and place itself on a least-loaded potential choice (ties are broken randomly). We now present the following theorem for the maximum load attained by $\mathcal{A}'(G,l)$.

Theorem E.1. *Suppose that G is a d-regular graph with girth at least $10l(\log\log n)^2$ and $d \in [3, \mathcal{O}(\log n)]$. Then, with high probability the maximum load attained by $\mathcal{A}'(G,l)$, denoted by m^*, is bounded from above as follows:*

I. If $\omega(1) \leqslant l \leqslant 4\gamma'_G$, where $\gamma'_G = \sqrt{\log_d n / r_G}$, then we have

$$m^* = \mathcal{O}\left(\frac{\log_d n \cdot \log\log n}{r_G l^2}\right).$$

II. If $l \geqslant 4\gamma'_G$, then we have

$$m^* = \mathcal{O}\left(\frac{\log\log n}{\log(l/\gamma_G)}\right).$$

The analysis of allocation algorithm $\mathcal{A}'(G,l)$ is almost the same as the algorithm for dense graphs so we only outline parts of the proof and notations that are slightly different. Let us start by defining an interference graph and show some of its properties (similar to Section 2)

Definition 5. *(Interference graph) For any given pair (G,l), interference graph $\mathcal{I}'(G,l)$ is a graph whose nodes are the set of potential choices in each $l \cdot r_G$-walk on G and two sets are connected if and only if they intersect each other.*

Lemma E.2. *We have that*

$$\begin{aligned}
|V(\mathcal{I}')| &= nd(d-1)^{l\cdot r_G - 1}/2, \\
\Delta(\mathcal{I}') &\leqslant (l+1)^2 d(d-1)^{l\cdot r_G - 1}.
\end{aligned}$$

Proof. Since G has girth at least $l \cdot r_G$, each $l \cdot r_G$-walk determines a unique set of potential choices and hence $|V(\mathcal{I}')|$ is the number of all lr_G-walks, that is,

$$|V(\mathcal{I}')| = nd(d-1)^{l\cdot r_G - 1}/2.$$

It is easy to show that, for every $0 \leqslant i \leqslant l$, each node can be the (ir_G)-th visited node of $d(d-1)^{lr_G-1}$ lr_G-walks. So every node is contained in at most $(l+1)d(d-1)^{lr_G-1}$ many walks. On the other hand, each set contains $l+1$ nodes and hence,

$$\Delta(\mathcal{I}') \leqslant (l+1)^2 d(d-1)^{l\cdot r_G - 1}.$$

□

Similar to Section 2, we can define the (λ, μ)-trees contained in $\mathcal{I}'$ and get the similar results in terms of $V(\mathcal{I}')$ and $\Delta(\mathcal{I}')$.

The shown results for dense graph including appearance probability of a c-loaded (λ, μ)-tree (Section 2) and witness graphs (Section 3) are based on the local properties of a regualr graph with degree at $\omega(\log n)$ and girth $10lh$. We also know that allocation algorithm $\mathcal{A}'(G, l)$ samples nodes from a neighborhood of radius $l \cdot r_G$. It only collects the load information every r_G-th visited node and ignore rest of the visited nodes. Now if we sort potential choices according to their distance from the starting node, say $(u_0, u_1 \cdots, u_l)$. It is easy to see that each u_i, $1 \leqslant i \leqslant l$, is chosen from a set of nodes of size at least $(d-1)^{r_G} = \omega(\log n)$, as the graph locally looks like a tree.

Roughly speaking, $\mathcal{A}'(G, l)$ reduces sparse d-regular graph G to a regular graph with degree $\omega(\log n)$ and hence the results also hold for $\mathcal{A}'$.

F A Lower Bound

In this section we derive a lower bound for the maximum load attained by allocation algorithms on dense and sparse graphs. Define

$$r_G := \begin{cases} 1 & \text{if } d = \omega(\log n)\ , \\ \lceil 2 \log_{d-1} \log n \rceil & \text{otherwise.} \end{cases}$$

Consider the generic balanced allocation $\mathcal{A}(G, l)$ that proceeds af follows: each ball takes a NBRW of length $l \cdot r_G$, and then it places itself on a least-loaded node among every r_G-th visited node. When $r_G = 1$, it means every visited node is considered as a potential choice. It is easy to see that the algorithm covers both classes of graph. Now we show the following lower bound for the maximum load.

Theorem F.1 (Lower Bound)**.** *Suppose that G be a d-regular n-vertex graph with girth at least $\omega(l)$, where $l \in [32r_G, \mathcal{O}(\gamma_G)]$ is an integer and $\gamma_G = \sqrt{\log_d n / r_G}$. Then with probability $1 - n^{-\Omega(1)}$ the maximum load attained by $\mathcal{A}(G, l)$ is at least*

$$\Omega\left(\frac{\log_d n}{r_G \cdot l^2}\right).$$

Proof. We know in each round the algorithm picks a random path of length $l \cdot r_G$. Let us define indicator random variable X_P for every path of length $l \cdot r_G$ as follows,

$$X_P := \begin{cases} 1 & \text{if } P \text{ is chosen at least } \tau \text{ times by } \mathcal{A}, \\ 0 & \text{otherwise,} \end{cases}$$

where τ will be specified later. The total number of paths of length $l \cdot r_G$, say s, is

$$nd(d-1)^{l \cdot r_G - 1}/2 \leqslant nd^{l \cdot r_G}/2.$$

Let P be an arbitrary path of length $l \cdot r_G$ in G. Thus we get

$$\begin{aligned} \mathbf{Pr}\left[X_P = 1\right] &= \sum_{i=\tau}^{n} \binom{n}{i} \left(\frac{1}{s}\right)^i \left(1 - \frac{1}{s}\right)^{n-i} \geqslant \left(\frac{n}{s \cdot \tau}\right)^{\tau} \left(1 - \frac{1}{s}\right)^{n} \\ &\geqslant \left(\frac{2}{d^{lr_G} \cdot \tau}\right)^{\tau} \left(1 - \frac{1}{s}\right)^{s} \geqslant d^{-(lr_G + \log_d \tau)\tau}/\mathrm{e}, \end{aligned} \tag{4}$$

where the second inequality follows from $n \leqslant s \leqslant n \cdot d^{lr_G}/2$. By setting

$$\tau = \frac{\log_d n}{6l \cdot r_G},$$

and using the fact that $\log_d \tau < \log_d \log_d n \leqslant r_G \leqslant l$ we get

$$(lr_G + \log_d \tau)\tau \leqslant \log_d n/6 + \log_d n/6 = \log_d n/3.$$

By substituting the above upper bound in (4), we get

$$\mathbf{Pr}\left[X_P = 1\right] = \Omega(n^{-1/3}).$$

Let us define the random variable $Y = \sum_{\text{all paths}} X_P$. By linearity of expectation we have

$$\mathbf{E}\left[Y\right] = s \cdot \mathbf{Pr}\left[X_P = 1\right] = (n \cdot d \cdot (d-1)^{lr_G - 1}/2)\Omega(n^{-1/3}) = \Omega(n^{2/3}). \tag{5}$$

It is easily seen that the random variables X_P and $X_{P'}$ are negatively correlated, which means for every P and P',

$$\mathbf{E}\left[X_P \cdot X_{P'}\right] \leqslant \mathbf{E}\left[X_P\right] \cdot \mathbf{E}\left[X_{P'}\right].$$

This implies that

$$\begin{aligned}\mathbf{Var}\left[Y\right] &= \sum_P (\mathbf{E}\left[X_P^2\right] - (\mathbf{E}\left[X_P\right])^2) + \underbrace{\sum_{P \neq P'} (\mathbf{E}\left[X_P X_{P'}\right] - \mathbf{E}\left[X_P\right]\mathbf{E}\left[X_{P'}\right])}_{\leqslant 0} \\ &\leqslant \sum_P \mathbf{E}\left[X_P^2\right] = \mathbf{E}\left[Y\right].\end{aligned}$$

Applying Chebychev's inequality and above inequality yield that

$$\mathbf{Pr}\left[Y = 0\right] \leqslant \mathbf{Pr}\left[|Y - \mathbf{E}\left[Y\right]| \geqslant \mathbf{E}\left[Y\right]\right] = \frac{\mathbf{Var}\left[Y\right]}{(\mathbf{E}\left[Y\right])^2} \leqslant \frac{1}{\mathbf{E}\left[Y\right]}.$$

By equality (5) we have that $\mathbf{E}\left[Y\right] = \Omega(n^{2/3})$. Therefore with probability at least $1 - \mathcal{O}(n^{-2/3})$ we have $Y \geqslant 1$, which means there exists a path P that is chosen at least τ times. Since every P contains $l+1$ choices, by the pigeonhole principle there is a node with load at least

$$\Omega\left(\frac{\tau}{l}\right) = \Omega\left(\frac{\log_d n}{r_G l^2}\right).$$

□